Atlas LSH Neural Networks: An Intuitive Overview

Modern neural networks typically decide how information flows through the network one layer at a time. Each layer influences the next, creating a chain of routing decisions as the computation proceeds.

Atlas LSH neural networks explore a different idea.

Instead of deciding the route through the network as computation unfolds, the network determines its route at the very beginning.

The central idea can be summarized as:

> For each input, the Parameter Atlas selects a sparse computational route through the network. Learning then sculpts the computations performed along that route.

The Parameter Atlas

Imagine planning a journey using a road atlas.

You first determine where you are and where you want to go. The atlas then suggests an appropriate route through the road network.

An Atlas neural network works in a similar way.

Given an input, such as an image or a sentence, the network first consults its **Parameter Atlas**. The atlas determines which parts of the network should participate in processing that particular input.

Only a subset of the network's parameters may be activated.

Different inputs may therefore follow different computational routes.

Nearby or similar inputs are expected to follow similar routes.

What is LSH?

LSH stands for **Locality Sensitive Hashing**.

A hash function usually converts data into a compact code. Traditional hash functions intentionally scramble information so that similar inputs produce unrelated outputs.

LSH does the opposite.

An LSH system attempts to preserve locality:

- similar inputs tend to produce similar codes;
- nearby inputs tend to activate nearby regions of the atlas;
- related inputs tend to share computational pathways.

One simple way to construct an LSH is to project the input onto many randomly chosen directions and record only whether each projection is positive or negative.

The resulting binary code acts as a compact description of the input's location within the atlas.

Sparse Computational Routes

Once the atlas has selected a route, only the parameters associated with that route become active.

For example:

Input image
↓
Atlas lookup
↓
Route selected
↓
Activate parameter subsets
↓
Compute output

This allows very large networks to operate sparsely.

The atlas effectively determines:

> Which parts of the network should work on this input?

Function Sculpting

After the route has been selected, learning takes place in the usual manner.

Gradient descent, backpropagation, Direct Feedback Alignment, or other learning algorithms can adjust the parameters that participated in the computation.

This learning process may be viewed as **Function Sculpting**.

The atlas determines where computation occurs.

Learning determines what computation occurs.

The network therefore consists of two conceptually separate systems:

1. **The Parameter Atlas**
    - selects routes;
    - activates parameters;
    - preserves computational locality.
2. **Function Sculpting**
    - learns useful computations;
    - adapts local functions;
    - improves task performance.

Atlas Networks and Mixture of Experts

The atlas concept naturally generalizes to mixture-of-experts architectures.

The atlas need not merely select individual parameters.

It could also select:

- entire subnetworks;
- specialist experts;
- groups of experts;
- different computational strategies.

For example:

Input
↓
Atlas
↓
Expert selection
↓
Specialized computation

Thus, expert selection can be viewed as another form of atlas-guided routing.

Biological Inspiration

One may also view biological brains through an atlas perspective.

Evolution appears to provide substantial innate structure:

- cortical organization;
- sensory maps;
- long-range connectivity patterns;
- specialized brain regions.

These structures define a relatively stable computational geography.

Experience then modifies synaptic strengths within this inherited structure.

From this viewpoint:

> Evolution builds the atlas. Experience sculpts the learned response.

The atlas need not encode behavior directly. Instead, it provides a computational landscape within which learning can efficiently adapt.

Summary

Atlas LSH neural networks explore a simple but potentially powerful idea:

> First decide where computation should occur. Then learn what computation should occur there.

The Parameter Atlas provides stable, locality-preserving computational routes, while Function Sculpting learns the computations performed along those routes.

Whether such architectures ultimately outperform conventional neural networks remains an open question, but they provide an intriguing alternative way of thinking about intelligence, learning, and sparse computation.